arXiv:quant-ph/0309152v2 6 Jun 2006

Blind quantum computation

Pablo Arrighi[1,*] and Louis Salvail[2,†]

[1]*Laboratoire Leibniz, Institut d'Informatique et de Mathématiques Appliquées de Grenoble (IMAG), CNRS UMR 5522, 46 Avenue Félix Viallet, 38031 Grenoble Cedex, France.*

[2]*BRICS, Department of Computer Science, University of Aarhus, Building 540, Ny Munkegade, Aarhus C-8000, Denmark.*

We investigate the possibility of *having someone carry out the work of executing a function for you, but without letting him learn anything about your input.* Say Alice wants Bob to compute some known function f upon her input x, but wants to prevent Bob from learning anything about x. The situation arises for instance if client Alice has limited computational resources in comparison with mistrusted server Bob, or if x is an inherently mobile piece of data. Could there be a protocol whereby Bob is forced to compute $f(x)$ *blindly*, i.e. without observing x? We provide such a blind computation protocol for the class of functions which admit an efficient procedure to generate random input-output pairs, e.g. factorization. The cheat-sensitive security achieved relies only upon quantum theory being true. The security analysis carried out assumes the eavesdropper performs individual attacks.

PACS numbers: 03.67.Dd
Keywords: Secure circuit evaluation, Secure two-party computation, Information hiding, Information gain versus disturbance, Quantum cryptography

I. INTRODUCTION

In the traditional secure two-party computation scenario[2, 11] Alice has secret input x, Bob has secret input y, and both of them wish to compute $f(x, y)$. The function f is of course well-known to the two parties; the usual example is that of two millionaires who wish to compare their wealth without disclosing how much they own[11]. Most protocols for secure two-party computation are symmetric with respect to the computing power each party should carry out during the execution. In these scenarios, if Alice knew Bob's input y she could compute $f(x, y)$ on her own without having to invest more computing power. Entering a secure two-party computation together with Bob will in general not help in diminishing Alice's computing power needed to evaluate f, and this is simply not the aim pursued. In fact, all implementations known by the authors of this paper require both Alice and Bob to invest more computing power than what is needed for the mere evaluation of f. For instance in [2] each gate performed by a party requires the other to perform the same gate, together with some extra encryption.

Unlike secure two-party computation, blind computation is fundamentally asymmetric. Alice is the only party with a secret input x, Bob is the only one able to compute f. Alice wants Bob to compute $f(x)$ without him learning *too much* about x. Thus an obvious motivation for Alice to enter a blind quantum computation together with Bob is to unload the computational task of computing f without having to compromise the privacy of her input. One could easily imagine this occurring in a Grid architecture, or in any client-server relation with a mistrusted server retaining the computational power. To make things more precise, suppose there were only a handful of fully operational large-scale quantum computers in the world, and some hungry academic decided to make use of her timeshare as scientist to crack some Swiss bank's RSA private key x. The hungry academic (Alice) will surely want to keep x secret from the authorities handling the quantum computer (Bob), so that she does not get suspected when subsequent international money transfers come to top up her meager income. But there may be other reasons to enter a blind computation protocol than mere computational power asymmetry. For instance Bob may possess some trapdoor information about the otherwise well-known function f. Or perhaps x may represent some mobile agent's code which ought to be protected against the malicious host upon which it runs. Others may see blind quantum computation as a somewhat philosophical issue: Is it possible to carry out some work for someone whilst being prevented from knowing what the work consists in?

In the classical setting, blind computation has first been studied by Feigenbaum [6]. It was shown that for some functions f, an instance x can be encrypted by $z = E_k(x)$ in such a way that Alice can recover $f(x)$ efficiently from k and $f(z)$. The construction cannot be extended easily to general classes of functions. In particular, blind computation of the discrete logarithm function (DLF) was shown possible but no blind computation of the RSA factoring function (FACF) is known. The infinite complexity hierarchy $P \subseteq NP = NP^{\emptyset} \subseteq NP^{NP} \subseteq \ldots \subseteq NP^{NP^{\cdot^{\cdot^{NP}}}} \subseteq \ldots$ (where NP^C stands for the class of language recognizable in non-deterministic polynomial

*Electronic address: pablo.arrighi@imag.fr
†Electronic address: salvail@brics.dk

time provided access to an oracle for problems of class C) is called the *polynomial-time hierarchy*. It is widely believed that every level in the polynomial-time hierarchy is strictly contained in the next one. However, proving or disproving this statement would be a major breakthrough in complexity theory. Abadi, Feigenbaum, and Kilian[1] have shown that no NP-hard problem can be computed blindly unless the polynomial-time hierarchy collapses at the third level. We conclude that it is very unlikely that any NP-hard problem can be computed blindly in the classical setting.

Even when computational assumptions are invoked [10], none of the currently known classical blind computation protocols applies to general classes of functions. Rather they take advantage of specific algebraic properties of particular functions. These constructions rely upon encryptions that are, in some sense, homomorphic with respect to function f. Clearly, very natural candidates for f are not known to have this property like for FACF. It is not surprising that such stringent requirements do not necessarily hold when Bob is running a quantum computer.

In this paper as in[1, 6], we are concerned with unconditional security, that is we do not make any computational assumptions upon eavesdropper Bob. Although we give Bob the opportunity to gain some Shannon information I about Alice's input x, we ensure that Bob's eavesdropping gets detected by Alice with a probability which rapidly increases with I. Any server Bob who wants to remain in business should clearly avoid such an a posteriori detection. Our goal consists of finding protocols for blind computation for which a good tradeoff between Bob's ability to avoid being detected and the amount of Shannon information about Alice's input can be established. Almost privacy was recently studied by Klauck [9] in a two-party computation setting which differs from the asymmetric scenario imposed by blind computations. Moreover, the security was only considered with respect to passive adversaries. We want our solution to apply to a wider class of functions than the one considered in the classical setting while being resistant to active adversaries. As far as we can tell, blind quantum computation has not been studied as such so far.

In section II we present the basic ideas of our blind quantum computation protocol, as well as the reasons which limit their use to a certain class of functions. In section III we review and adapt a recent result in the Information versus Disturbance tradeoff literature. In section IV we formalize the protocol and give a proof of its security. We conclude in section V and mention possible extensions of this work.

II. PRINCIPLES OF A SOLUTION

Let us now explain the basic principles underlying our blind quantum computation protocol. Suppose Alice wants Bob to compute $f(x)$ whilst keeping x secret. Moreover suppose Bob possesses a quantum computer which implements f, i.e. he is able to implement a unitary transform U such that $U|q\rangle|0\rangle = |q\rangle|f(q)\rangle$ for all inputs, q. In order to achieve her purpose Alice could hide her true input $|x\rangle$ amongst superpositions of other potential inputs $\frac{|q\rangle+i|q'\rangle}{\sqrt{2}}$ (which we later refer to as 'quantum decoys') and send all this to Bob for him to execute U. Now if Bob attempts a measure so as to determine $|x\rangle$ he will run the risk of collapsing the superpositions. Alice may detect such a tampering when she retrieves her results. The above suggestion has a weakness however: Alice is not returned $\frac{|q\rangle+i|q'\rangle}{\sqrt{2}}$, but

$$U\frac{|q\rangle + i|q'\rangle}{\sqrt{2}}|0\rangle = \frac{|q; f(q)\rangle + i|q'; f(q')\rangle}{\sqrt{2}},$$

the result of Bob's computation upon the superposition Alice had sent. Since Alice does not want to compute f herself she is in general unable to check upon the integrity of such states. To get an intuition of why this is consider the effects of tracing out the result register whenever $f(q)$ is different from $f(q')$.

$$\begin{aligned}
&\mathrm{Tr}_2\Big(\frac{|q; f(q)\rangle + i|q'; f(q')\rangle\langle q; f(q)| - i\langle q'; f(q')|}{2}\Big)\\
&= \big(|q\rangle\langle q|\mathrm{Tr}(|f(q)\rangle\langle f(q)|) - i|q\rangle\langle q'|\mathrm{Tr}(|f(q)\rangle\langle f(q')|)\\
&+ i|q'\rangle\langle q|\mathrm{Tr}(|f(q')\rangle\langle f(q)|) + |q'\rangle\langle q'|\mathrm{Tr}(|f(q')\rangle\langle f(q')|)\big)/2\\
&= \frac{|q\rangle\langle q| + |q'\rangle\langle q'|}{2}
\end{aligned}$$

In other words once such a trace-out has been performed the state is either $|q\rangle$ with 0.5 probability, or $|q'\rangle$ with 0.5 probability, i.e. it makes no difference whether Bob performed a measurement in the computational basis or not.

There are many computational problems, however, for which this obstacle can be circumvened. For example say f takes composite numbers into the list of their integer factors. Then Alice can easily (at the cost of a few multiplications) prepare several input-output pairs $\{(q, f(q))\}$. Thus if Alice hides her true input $|x\rangle$ amongst superpositions $\frac{|q\rangle+i|q'\rangle}{\sqrt{2}}$ generated in this manner, she will later be able to check whether $\frac{|q;f(q)\rangle+i|q';f(q')\rangle}{\sqrt{2}}$ are indeed being returned. Formally the idealized class of functions for which our protocol will work is defined as follows:

Definition 1 (Random verifiable functions) *Let S and S' denote two finite sets. A function $f : S \to S'$ is* random verifiable *if and only if there exists, for all N, an efficient probabilistic process which generates N input-output pairs $\{(q, f(q))\}$ and such that the inputs (the q's) are uniformly distributed in S.*

There are several promised problems for which we can define functions that are random verifiable. Consider the language *RSA*-composite which contains natural numbers of a fixed size that can be expressed by the product

of two primes of the same size. The function f that returns the prime factors is also random verifiable. In this case, f can be computed efficiently on a quantum computer but not, as far as we know, on a classical computer. Another example can be obtained from the graph isomorphism problem. Let $L_{e,v}$ be the set of all pairs of isomorphic graphs with e edges and v vertices. We define function $f : L_{e,v} \mapsto S_e$, where S_e is the set of all permutations among v elements, as $f(G_0, G_1) = \sigma$ such that $\sigma(G_0) = G_1$. It is easy to verify that f is random verifiable. The following efficient classical computation does the job:

- Pick a random permutation $\sigma \in S_e$,
- Generate a random graph G_0 with e edges and v vertices,
- Output $((G_0, \sigma(G_0)), \sigma)$.

Although f is random verifiable by an efficient classical algorithm, it is not known whether even a quantum computer can evaluate f efficiently.

In this paper, we provide a blind quantum computation protocol for random verifiable functions together with a thorough security analysis. The cheat-sensitive security achieved relies upon the laws of physics only. It is expressed using the vocabulary of information theory. As was hinted in this section our analysis will crucially depend upon the tradeoff between Bob's information gain about Alice's true input (a canonical basis state) and the disturbance he induces upon superpositions of potential inputs (pairwise superpositions of canonical basis states).

III. INFORMATION GAIN VERSUS DISTURBANCE TRADEOFF

Say Alice draws out a state from an ensemble of quantum states, sends it to Bob, and later retrieves it. How much information can Bob learn about the state, and what, then, is the probability that Alice can detect Bob's eavesdropping? Questions of Information Gain versus Disturbance tradeoff were first investigated by Fuchs and Peres [7], who considered a seemingly simple scenario involving two equiprobable non-orthogonal pure states. But the formula they obtained is relatively complex and the methods employed are somewhat difficult to export to our setting. In order to construct a blind quantum computation protocol we needed to quantify the disturbance upon pairwise superpositions of n-dimensional canonical basis states, as induced when Bob seeks to learn information about the canonical basis. A tradeoff formula for this problem was given in [3]. Proposition 1 rephrases this result in terms of induced fidelity and letting Bob and Alice be the same person.

Scenario 1 (One quantum decoy) *Consider a quantum channel for transmitting n-dimensional systems having canonical orthonormal basis $\{|j\rangle\}$.*
Suppose Alice's message words are drawn out of the canonical ensemble $\{(1/n, |j\rangle)\}_{j=1..n}$, whilst her quantum decoys *are drawn out of the pairing ensemble $\{(1/n^2, \frac{|j\rangle+i|k\rangle}{\sqrt{2}})\}$. Alice sends, over the quantum channel, either a message word or a decoy, which she later retrieves.*
Whenever she sends a quantum decoy $\frac{|j\rangle+i|k\rangle}{\sqrt{2}}$ she later measures the retrieved system with $\{P_{intact} = \left(\frac{|j\rangle+i|k\rangle}{\sqrt{2}}\right)\left(\frac{\langle j|-i\langle k|}{\sqrt{2}}\right),\ P_{tamper} = \mathbb{I} - P_{intact}\}$ so as to check for tampering.
Suppose Bob is eavesdropping the quantum channel, and has an interest in determining Alice's message words.

FIG. 1: One quantum decoy.

$A:$ Draw t in $T = \{(p, go), (1-p, nogo)\}$
$A:$ If $t = go$ draw $s = m$ in $M = \{(1/n, |j\rangle)\}_{j=1..n}$
$A:$ If $t = nogo$ draw $s = d$ in $D = \{(1/n^2, \frac{|j\rangle + i|k\rangle}{\sqrt{2}})\}$
$A:\ s \longrightarrow B$
$B:$ Draw x, s' in $X, S = \{(||M_x s||^2, |x\rangle \otimes M_x s/||M_x s||\}_x$ with $\{M_x\}$ a generalized measurement.
$B:\ s' \longrightarrow A$
$A:$ If $t = go$ draw y in $Y = \{(0, tamp), (1, notamp)\}$
$A:$ If $t = nogo$ draw y in $Y = \{(||P_{tamper} s'||^2, tamp), (||P_{intact} s'||^2, notamp)\}$ with $P_{tamper} = \mathbb{I} - P_{intact},\ P_{intact} = ss^\dagger$.

Proposition 1 (One quantum decoy)
Referring to Scenario 1 and its formalization in Figure 1, suppose Bob performs an attack such that, whenever a message word gets sent, he is able to identify which with probability G (mean estimation fidelity).
Then, whenever a quantum decoy gets sent, the probability F (induced fidelity) of Bob's tampering not being detected by Alice is bounded above under the following tight inequality:

$$F \leq \frac{1}{2} + \frac{1}{2n}\left(\sqrt{G} + \sqrt{(n-1)(1-G)}\right)^2 \qquad (1)$$

For optimal attacks G varies from $\frac{1}{n}$ to 1 as F varies from 1 to $\frac{1}{2}+\frac{1}{2n}$.

Now imagine that Scenario 1 gets repeated N times round, and that Alice happens to send only decoys.

Scenario 2 (N Quantum decoys)
Step 0. Alice prepares a pool of $N+1$ quantum states consisting of one message word together with N quantum decoys.
Step 1. Alice sends Bob one quantum state drawn at

random amongst those remaining in the pool.
Step 2. Alice awaits to retrieve the quantum state she sent.
Step 3. If Alice sent a quantum decoy she measures the retrieved system so as to check for tampering.
Step 4. If the pool is empty Alice stops the protocol, else she proceeds again with Step 1.
Step 5. Alice publicly announces the position p at which she sent her message word.

This scenario is formalized in Figure 2.

FIG. 2: N QUANTUM DECOYS / INDIVIDUAL ATTACKS

A : draw p in $P = \{(1/(N+1), i), \}_{i=0\ldots N}$
For $r = 0 \ldots N$:
A : If $r \neq p$ draw t in $T = \{(1, go), (0, nogo)\}$
else draw t in $T = \{(0, go), (1, nogo)\}$
A : If $t = go$ draw $s = m$ in $M = \{(1/n, |j\rangle)\}_{j=1..n}$
A : If $t = nogo$ draw $s = d$ in
$D = \{(1/n^2, \frac{|j\rangle + i|k\rangle}{\sqrt{2}})\}$
A : $s \longrightarrow B$
B : Draw x_r, s' in
$X, S = \{(||M_x^{(r)} s||^2, |x\rangle \otimes M_x^{(r)} s / ||M_x^{(r)} s||\}_r$
with, for all r, $\{M_x^{(r)}\}$
a gen. mesurement upon a bipartite system.
B : $s' \longrightarrow A$
A : If $t = go$ draw y_r in
$Y = \{(0, tamp), (1, notamp)\}$
A : If $t = nogo$ draw y_r in
$Y = \{(||P_{tamper} s'||^2, tamp), (||P_{intact} v||^2, notamp)\}$
with $P_{tamper} = \mathbb{I} - P_{intact}, P_{intact} = ss^\dagger$.
$A : p \longrightarrow B$
$B : guess = f_p(x_0, \ldots, x_N)$

Corollary 1 (N quantum decoys/individual attacks) *Referring to Referring to Scenario 2, suppose Bob only performs individual attacks, i.e. independent of each other at each round, as formalized in Figure 2.*
The probability of Bob reaching round m ($0 \leq m \leq N$) without being caught tampering is bounded above under the following tight inequality:

$$p(Bob\ reaches\ m) \leq \prod_{i=1, i\neq p}^{m-1} F(G_i) \qquad (2)$$

where G_i stands for Bob's mean estimation fidelity, if p is announced equal to i, about the message word sent at round p.

Proof. Within the for loop of Figure 2 the scenario which gets repeated is exactly that of Figure 1, for which Proposition 1 applies independently at each round. □

In the above scenario Bob's attacks are somewhat memoryless. Bob's measurements do not depend upon previous outcomes nor upon any ancilla quantum system which he might keep throughout the protocol. This is what enables us to apply Proposition 1 at the level of each individual transmission [3], i.e. to assume that that the probabilities $\{p(\text{Bob passing round } i) = F(G_i)\}_{i=1..N}$ are independent from each other and hence that that Bob's chances of not being detected at all are bounded by $\prod_{i=1..N} F(G_i)$.
Now say Bob was to keep an ancillary quantum system entangled with a quantum decoy sent at a previous round, and then perform a coherent quantum measurement upon another quantum decoy and the ancillary quantum system at a later round – could this correlate his probabilities of getting caught in a favorable manner? We argue that it is not so in the following conjecture, by making use of a standard argument. Formal proofs of probabilistic security protocols are known to be an extremely delicate matter requiring delicate notions of process equivalences. In quantum information theory such rigorous frameworks have not yet appeared and seem to be needed here – we will only provide the reader with a number of intuitions which strongly support our statement.

Conjecture 1 (N quantum decoys/coherent attacks) *Referring to Scenario 2, suppose Bob performs general attacks, i.e. which may depend from each other at every round, as formalized in Figure 3.*
The probability of Bob reaching round m ($0 \leq m \leq N$) without being caught tampering is bounded above under the following tight inequality:

$$p(Bob\ reaches\ m) \leq \prod_{i=1, i\neq p}^{m-1} F(G_i)$$

where G_i stands for Bob's mean estimation fidelity, if p is announced equal to i, about the message word sent at round p.

The following arguments support our claim. Note that in this Figure 3 we allow Bob to perform the most general attack possible: his generalized measurements $\{M_x^{(r)}\}$ depend upon the round r; they may entangle the ancillary quantum system a to the state sent by Alice s for later use (thus the systems a' and s' may be entangled); they may depend upon previous measurement outcomes via the contents of the ancillary quantum system a; or they could keep a entangled but unmeasured until the final round provided that for $r < N$ the statistics of $\{M_x^{(r)}\}$ do not depend on a. We now reason by contradiction.
Suppose $p(\text{Bob reaches } m) > \prod_{i=1, i\neq p}^{m-1} F(G_i)$. Then

FIG. 3: N QUANTUM DECOYS / COHERENT ATTACKS

A : draw p in $P = \{(1/(N+1), i), \}_{i=0\ldots N}$
$B : a = \psi$
For $r = 0 \ldots N$:
  A : If $r \neq p$ draw t in $T = \{(1, go), (0, nogo)\}$
    else draw t in $T = \{(0, go), (1, nogo)\}$
  A : If $t = go$ draw $s = m$ in $M = \{(1/n, |j\rangle)\}_{j=1..n}$
  A : If $t = nogo$ draw $s = d$ in
    $D = \{(1/n^2, \frac{|j\rangle + i|k\rangle}{\sqrt{2}})\}$
  A : $s \longrightarrow B$
  B : Draw x_r, s', a' in
    $X, S, A = \{(||M_x^{(r)} sa||^2, |x\rangle \otimes M_x^{(r)} sa/||M_x^{(r)} sa||\}_r$
    with, for all r, $\{M_x^{(r)}\}$
    a gen. mesurement upon a bipartite system.
  B : $s' \longrightarrow A$
  A : If $t = go$ draw y_r in
    $Y = \{(0, tamp), (1, notamp)\}$
  A : If $t = nogo$ draw y_r in
    $Y = \{(||P_{tamper} s'||^2, tamp), (||P_{intact} v||^2, notamp)\}$
    with $P_{tamper} = \mathbb{I} - P_{intact}, P_{intact} = ss^\dagger$.
  B : $a = a'$
$A : p \longrightarrow B$
$B : guess = f_p(x_0, \ldots, x_N)$

there exists a k for which

$$p(\text{Bob reaches } k) \leq \prod_{i=1, i\neq p}^{k-1} F(G_i) \quad \text{and}$$

$$p(\text{Bob reaches } k{+}1) > \prod_{i=1, i\neq p}^{k} F(G_i).$$

For such a k we thus have

$$p(\text{Bob reaches } k{+}1 | \text{Bob reaches } k) > F(G_k). \qquad (3)$$

In other words Bob, on the k^{th} round, due to the state of the ancillary system a at this round, is capable of collecting mean estimation G_k about a message word whilst remaining undetected with probability more than $F(G_k)$ upon a quantum decoy. However a is absolutely uncorrelated with s for our purpose, because:

- the quantum decoys and the message words are undistinguishable since $(1/n^2)\sum_{jk}(|j\rangle + i|k\rangle)(\langle j| - i\langle k|) = (1/n)\sum_i |i\rangle\langle i|$. Hence a cannot hold any information about whether the message word is sent at round k;
- the quantum decoys are picked up independently from one another and independently from the message words, hence if the message word is sent at round k, a does not hold any complementary information about the message word and does not modify G_k;
- the quantum decoys are picked up independently from one another and independently from the message words, hence if a quantum decoy is sent at round k, a does not hold any complementary information about the subspace of the quantum decoy which needs to be preserved and hence does not modify $F(G_k)$.

In other words Bob could have, for the purpose of optimizing his information gain versus disturbance tradeoff at round k, come up with just as good an a by playing the first $k-1$ rounds of the protocol with Charlie instead. Hence the situation at round k is in contradiction with Proposition 1. ⊠

The next section also makes use of the following mathematical result, whose direct proof was shown to us by Prof. Frank Kelly.

Lemma 1 (Concavity of circular products) *Consider $f : [0,1] \to [0,1]$ a concave, continuous function and $\{x_i\}_{i=1\ldots N+1}$ a set of real numbers in the interval $[0,1]$.*
Suppose the sum $t = \sum_{i=1}^{N+1} x_i$ is fixed. We have

$$\frac{1}{N+1}\sum_{p=1}^{N+1}\Big(\prod_{i=1, i\neq p}^{i=N+1} f(x_i)\Big) \leq f\big(\frac{t}{N+1}\big)^N.$$

Proof. By definition of concavity one has

$$\frac{1}{2}\big(f(x_1) + f(x_2)\big) \leq f\big(\frac{x_1+x_2}{2}\big) \qquad (4)$$

$$\text{and} \quad f(x_1)f(x_2) \leq f\big(\frac{x_1+x_2}{2}\big)^2, \qquad (5)$$

where the latter equation trivially derives from $f(x_1)f(x_2) \leq \big(\frac{f(x_1)+f(x_2)}{2}\big)^2$. Let us now show that

$$\frac{1}{N+1}\sum_{p=1}^{N+1}\prod_{i=1, i\neq p}^{i=N+1} f(x_i) \leq \frac{1}{N+1}\sum_{p=1}^{N+1}\prod_{i=1, i\neq p}^{i=N+1} f(y_i), \quad (6)$$

where $y_1 = y_2 = \frac{x_1+x_2}{2}$ and $y_i = x_i$ for $i = 3 \ldots N+1$. This result is in fact obtained by combining (summing) two inequalities:

$$\big(f(x_1) + f(x_2)\big)\prod_{i=3}^{N+1} f(x_i) \leq \big(f(y_1) + f(y_2)\big)\prod_{i=3}^{N+1} f(y_i)$$

$$f(x_1)f(x_2)\sum_{p=3}^{N+1}\prod_{i=3, i\neq p}^{i=N+1} f(x_i) \leq f(y_1)f(y_2)\sum_{p=3}^{N+1}\prod_{i=3, i\neq p}^{i=N+1} f(y_i)$$

where former stems from Equation (4) and $f(x)$ positive, whilst the latter stems from Equation (5) and $f(x)$ positive.

Equation (6) expresses the fact that, whenever two elements x_i and x_j, $i \neq j$ are replaced by their mean, the value of

$$\pi(\underline{x}) \equiv \frac{1}{N+1} \sum_{p=1}^{N+1} \Big(\prod_{i=1, i\neq p}^{i=N+1} f(x_i) \Big)$$

is increased. Now let us define $\{\underline{x}^{(k)}\}$ a sequence of vectors such that $\underline{x}^{(1)} = (x_1, x_2, \ldots, x_{N+1})$, and $\underline{x}^{(k)}$ is formed from $\underline{x}^{(k-1)}$ by replacing both the largest and the smallest component by their mean. As k goes to infinity this sequence of vectors tends to $\underline{x}^{(\infty)} = (\frac{t}{N+1}, \frac{t}{N+1}, \ldots)$. By Equation (6) we have $\{\pi(\underline{x}^{(k)})\}$ an increasing sequence of real numbers. As k goes to infinity, and since $\pi(\underline{x})$ is continuous in $\underline{x}$, this sequence of real numbers tends to

$$\pi(\underline{x}^{(\infty)}) = f\big(\frac{t}{N+1}\big)^N.$$

This limit must therefore provide, for all $\underline{x}$ having components summing to t, a tight upper bound on the value of $\pi(\underline{x})$. □

IV. PROTOCOL AND SECURITY

We are now set to give our blind quantum computation protocol:

Protocol 1 (Interactive version) *Alice wants Bob to compute $f(x)$ whilst keeping her input x secret. Here f designates a random verifiable function implemented on a quantum computer by a unitary evolution U.*

Step 0. Alice efficiently computes $2N$ random input-solution pairs $(q, f(q))$ and prepares a pool of $N+1$ quantum states consisting of her true input $|x\rangle$ together with N quantum decoys $\frac{|q\rangle + i|q'\rangle}{\sqrt{2}}$.

Step 1. Alice sends Bob one quantum state $|\psi\rangle$ drawn at random amongst those remaining in the pool.

Step 2. Bob supposedly computes $U|\psi\rangle|0\rangle$ and sends the result back to Alice.

Step 3. If $|\psi\rangle$ was a quantum decoy $\frac{|q\rangle + i|q'\rangle}{\sqrt{2}}$ Alice measures the retrieved system with

$$\big\{ P_{intact} = \frac{1}{2}\big(|q\, f(q)\rangle + i|q' f(q')\rangle\big)\big(\langle q\, f(q)| - i\langle q' f(q')|\big)$$
$$P_{tamper} = \mathbb{I} - P_{intact} \big\},$$

so as to check for tampering. tampering she stops. If on the other hand $|\psi\rangle$ was her true input Alice reads off $f(x)$.

Step 4. If the pool is empty Alice stops the protocol, else she proceeds again with Step 1.

Quantum theory is helpful for detecting observation by a mistrusted party through the induced disturbance. For this reason quantum cryptography has seen the rise of *cheat-sensitive* protocols where 'Either party may be able to evade the intended constraints on information transfer by deviating from these protocols. However, if they do, there is a non-zero probability that the other will detect their cheating' [8]. When the probability of detecting the cheating is one, the protocol may also be referred to as *cheat-evident* [5].

The security of our protocol is cheat-sensitive, as is rigorously described and quantified in the following claim. The security of our protocol may also be referred to as cheat-evident, in the sense that Alice's detection probability tends to 1 in the limit where N tends to infinity. Moreover for a fixed information gain by Bob, Alice's detection probability approaches 1 exponentially with N.

Claim 1 (Statement of security) *Referring to Protocol 1 suppose Bob has no a priori information about Alice's true input x. Let $I \in [0, \log(n)]$ be Bob's mutual information about Alice's true input x at the end of the protocol. Let $D \in [0, (1/2)^N]$ be the probability of Alice detecting Bob's tampering. Provided that Bob makes only individual attacks, the protocol ensures that $\forall G \in [\frac{1}{n}, 1]$,*

$$\big[I = \log(n) + \log(G) \;\; \Rightarrow \;\; D \geq 1 - F(G)^N \big].$$

Hence we have equivalently

$$D \geq 1 - F(2^{I - \log(n)})^N.$$

Proof.

We prove that the claim holds for a weakened form of Protocol 1, where we add:

Step 5. Alice publicly announces the position in which she sent her true input $|x\rangle$.

Until this stage, however, Bob has no means of knowing at which round true input $|x\rangle$ was sent. This is because we have assumed he has no a priori knowledge about the true input. In his view the state was drawn from the canonical ensemble $\{(1/n, |j\rangle)\}_{j=1..n}$, whilst the quantum decoys were drawn from the pairing ensemble $\{(1/n^2, \frac{|j\rangle + i|k\rangle}{\sqrt{2}})\}$, but the two are undistinguishable for they both have density matrix $\mathbb{I}/n$. We are, therefore, in the precise case of Corollary 1. Without loss of generality we can assume Bob's attack yields him mean estimation fidelity G_i about Alice's true input whenever the position is later announced equal to i. Let $\overline{G} = \sum_p G_p/(N+1)$.

First we prove that $[I = \log(n) + \log(G) \Rightarrow \overline{G} \geq G]$.

Say the true input is at position p. In this situation Bob's best chance of guessing the true input is G_p (by definition) and thus his Shannon uncertainty H_p about Alice's true input is bounded as follows

$$\begin{aligned} H_p &\equiv \sum -p(x|\text{Bob's outcome}) \log(p(x|\text{Bob's outcome})) \\ &\geq -\lfloor \frac{1}{G_p} \rfloor G_p \log(G_p) - (1 - \lfloor \frac{1}{G_p} \rfloor G_p) \log(1 - \lfloor \frac{1}{G_p} \rfloor G_p) \\ &\geq -\log(G_p). \end{aligned}$$

The RHS of the last line is often referred to as the 'min-entropy' sometimes denoted H_∞ and is commonly used to bound uncertainties in the above manner (i.e. Shannon uncertainty is always at least H_∞). As a consequence Bob's mutual information I_p satisfies

$$I_p \leq \log(n) + \log(G_p).$$

Averaging over all possible positions $p = 1 \ldots N+1$ Bob's mutual information satisfies

$$\begin{aligned} I &= \sum_{p=1}^{N+1} \frac{1}{N+1} I_p \\ &\leq \log(n) + \sum_{p=1}^{N+1} \frac{1}{N+1} \log(G_p) \\ &\leq \log(n) + \log(\overline{G}) \end{aligned}$$

where the third line was obtained using the concavity of $x \mapsto \log(x)$. Hence we have

$$I = \log(n) + \log(G) \leq \log(n) + \log(\overline{G}).$$

Since $x \mapsto \log(x)$ is crescent we conclude that $G \leq \overline{G}$. Second we prove that $[\overline{G} \geq G \Rightarrow D \geq 1 - F\big(G\big)^N]$. Since we have assumed individual attacks Corollary 1 applies, and so Bob is undetected with probability

$$p(\text{undetected}|p) \leq \prod_{i=1, i\neq p}^{N+1} F(G_i).$$

Let us now average the above over all possible positions $p = 1 \ldots N+1$. The probability that Bob's tampering remains undetected by Alice satisfies

$$\begin{aligned} p(\text{undetected}) &= \frac{1}{N+1} \sum_{p=1}^{N+1} p(\text{undetected}|p) \\ &\leq \frac{1}{N+1} \sum_{p=1}^{N+1} \Big(\prod_{i=1, i\neq p}^{N+1} F(G_i) \Big) \\ &\leq F\big(\overline{G}\big)^N \\ D &\geq 1 - F\big(\overline{G}\big)^N. \end{aligned}$$

where the third line was obtained using Lemma 1 upon the concave, continuous function $x \mapsto F(x)$. Since $x \mapsto 1 - F\big(x\big)^N$ is crescent and $\overline{G} \geq G$ we conclude that

$$D \geq 1 - F\big(\overline{G}\big)^N \geq 1 - F\big(G\big)^N.$$

□

Protocol 1 requires $N+1$ communications between Alice and Bob. One could suggest a modification whereby Alice would send Bob her whole pool (as prepared in *Step 0*), and later proceed to check upon the integrity of each element of the pool which Bob returns, apart from her true input. Formally this yields the following protocol:

Protocol 2 (Non-interactive version) *Alice wants Bob to compute $f(x)$ whilst keeping her input x secret. Here f designates a random verifiable function implemented on a quantum computer by a unitary evolution U.*

Step 0. Alice efficiently computes $2N$ random input-solution pairs $(q, f(q))$ and prepares a pool of $N+1$ quantum states consisting of her true input $|x\rangle$ together with N quantum decoys $\frac{|q\rangle + i|q'\rangle}{\sqrt{2}}$.

Step 1. Alice sends Bob the large quantum state $\bigotimes_{i=1}^{N+1} |\psi_i\rangle$ consisting of a random permutation of all elements of the pool.

Step 2. Bob supposedly computes $\bigotimes_{i=1}^{N+1} U|\psi_i\rangle|0\rangle$ and sends the result back to Alice.

Step 3. For each location i, if $|\psi_i\rangle$ was a quantum decoy $\frac{|q\rangle + i|q'\rangle}{\sqrt{2}}$ Alice measures

$$\Big\{ P_{intact} = \frac{1}{2}\big(|q; f(q)\rangle + i|q'; f(q')\rangle\big)\big(\langle q; f(q)| - i\langle q'; f(q')|\big)$$
$$P_{tamper} = \mathbb{I} - P_{intact} \Big\}.$$

so as to check for tampering. If on the other hand $|\psi_i\rangle$ was her true input Alice reads off $f(x)$.

When Bob is restricted to individual attacks (non-coherent attacks, i.e. Bob measures each quantum state in the pool individually) then Claim 1 holds also for Protocol 2. We omit the proof of this since it is similar, and in fact simpler than the one given for Protocol 1. Now suppose Conjecture 1 was verified. This would immediately entail that Claim 1 holds also for coherent attacks for Protocol 1. Hence we believe that Protocol 1 can resist the most general attack. A thought-provoking question is whether this is still the case of Protocol 2. Is it the case that interactivity contributes, to some extent, to a limitation of Bob's possible attacks?

V. CONCLUDING REMARKS

We have investigated the possibility of *having someone else carrying out the evaluation of a function for you without letting him learn anything about your input.* We gave a blind computation protocol for the class of functions which admit an efficient procedure to generate random input-output pairs. The protocol relies upon quantum physical information gain versus disturbance tradeoffs [3] to achieve cheat-sensitive security against individual attacks: whenever the server gathers $\log(n) + \log(G)$ bits of Shannon information about the input, he must get caught with probability at least $1 - F(G)^N$ (where n denotes the size of the input and N is a security parameter). Moreover the server cannot distinguish a weary client who uses the blind computation protocol (sending one true input amongst N decoys) from a normal client who simply makes repeated use of the server (sending $N+1$ true inputs). Thus if the server wanted to deny

his services to suspected users of the protocol, he would also have to refuse the normal clients. We have conjectured that the same security properties hold for general, coherent attacks.

Our protocol could be improved in several directions. In terms of costs one may hope to reduce the set of quantum gates needed by Alice to prepare her transmissions [4]; lower the size of the transmissions; lower the number of rounds required. We leave it as an open problem to find the security properties of the non-interactive version of our protocol when Bob is allowed coherent attacks.
In terms of functionality one may wish to achieve tamper prevention (preventing Bob from learning about x ever) rather than tamper detection (preventing Bob from learning about x without being detected, i.e. cheat-sensitiveness). Protocol 1 provides the latter to some degree, since its interactivity allows Alice to avoid sending her true input x whenever she detects tampering upon her quantum decoys in the previous rounds. However we have not provided an analysis for a tamper-prevention-like security property. Another challenge would be to extend/identify the class of functions admitting a blind quantum computation protocol. This may have consequences in quantum complexity theory, as was the case in the classical setting [1]. For instance if one was to prove that the blind quantum computation protocol had no interest as a secure way of discharging Alice computationally - because all the random verifiable functions turn out to be easy to perform on a quantum computer - then random verifiability would impose itself as an elegant property for the quantum polynomial class.

Acknowledgments

P.J.A would like to thank Prof. Frank Kelly for his self-contained proof of Lemma 1, Dr. Anuj Dawar for proof-reading, EPSRC, Marconi, the Cambridge European and Isaac Newton Trusts, and the European Union Marie Curie Fellowship scheme for financial support. L. S. would like to thank Ivan Damgård for enlightening discussions and the EU project PROSECCO for financial support. Both would like to thank the anonymous referee for important comments.

[1] M. Abadi, J. Feigenbaum, J. Kilian, *On Hiding Information from an Oracle*, Journal of Computer and System Sciences, **39**(1), 21-50, (1989).
[2] M. Abadi, J. Feigenbaum, *Secure Circuit Evaluation*, Journal of Cryptology, **2**(1), 1-12, (1990).
[3] P. Arrighi, *Quantum Decoys*, Int. J. of Quantum Information, **2**(3), 341-351, (2004).
[4] A.M. Childs, *Secure assisted quantum computation*, tech. report MIT-CTP 3211, arXiv:quant-ph/0111046.
[5] R. Colbeck, A. Kent, *Variable bias coin tossing: conspiring with fate to save the fate*, arXiv:quant-ph/0508149.
[6] J. Feigenbaum, *Encrypting Problem Instances, or, ..., Can You Take Advantage of Someone Without Having To Trust Him?*, Proc. CRYPTO'85, 477-488, Springer, (1986).
[7] C. Fuchs, *Information Gain vs. State Disturbance in Quantum Theory*, arXiv:quant-ph/9611010 and C. Fuchs, A. Peres, arXiv:quant-ph/9512023.
[8] L. Hardy, A. Kent, *Cheat sensitive quantum bit commitment*, Phys. Rev. Lett, **92**, 157901, (2004).
[9] H. Klauck, *On Quantum and Approximate Privacy*, Proc. of the 19^{th} Annual Symposium on Theoretical Aspects of Computer Science, LNCS, **2285**, 335-365, Springer, (2002).
[10] T. Sander, C.F. Tschudin *Protecting Mobile Agents Against Malicious Hosts*, Mobile Agents and Security, LNCS, **1419**, 44-61, Springer, (1998).
[11] A.C. Yao, *How to Generate and Exchange Secrets*, Proc. of the 27^{th} Annual Symposium on Foundations of Computer Science, 162-167, IEEE Computer Society Press, (1986).